



US006608812B1

(12) **United States Patent**
Wils et al.

(10) **Patent No.:** US 6,608,812 B1
(45) **Date of Patent:** Aug. 19, 2003

(54) **TRUNK ARRANGEMENT IN A COMPUTER NETWORK WITH INTRA TRUNK SWITCH CONNECTIONS**

(75) Inventors: **Joris Johannes Maria Wils**, Acton, MA (US); **James Scott Hiscock**, Rockport, MA (US)

(73) Assignee: **3Com Corporation**, Santa Clara, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/232,105**

(22) Filed: **Jan. 15, 1999**

(51) **Int. Cl.**⁷ **G01R 31/08**; G06F 11/00; G08C 15/00; H04J 1/16; H04J 3/14; H04L 1/00; H04L 12/26; H04L 12/28; H04L 12/56

(52) **U.S. Cl.** **370/217**; 370/389

(58) **Field of Search** 370/216, 217, 370/225-228, 389, 400

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,140,585 A * 8/1992 Tomikawa 370/354

5,398,236 A * 3/1995 Hemmady et al. 370/218
5,473,598 A * 12/1995 Takatori et al. 370/399
6,058,116 A * 5/2000 Hiscock et al. 370/401
6,088,330 A * 7/2000 Bruck et al. 370/228
6,111,852 A * 8/2000 Leung et al. 381/101
6,195,351 B1 * 2/2001 Hiscock et al. 370/389

* cited by examiner

Primary Examiner—Dang Ton

Assistant Examiner—Kevin C. Harper

(74) *Attorney, Agent, or Firm*—McGlew and Tuttle, P.C.

(57) **ABSTRACT**

A trunk switch of a trunk cluster for connecting a plurality of edge devices to each other, where each edge device can communicate with all the other edge devices through all of the trunk switches. Each trunk switch contains an intra cluster port and intra cluster links **35** connect the intra cluster ports of different trunk switches. When an edge device becomes disconnected from one trunk switch, but is still connected to another trunk switch, the trunk switch from which it has become disconnected transfers the data destined for the disconnected edge device to the intra cluster port which then transfers that data to another trunk switch. The another trunk switch, upon receiving the intra cluster data, then forwards that data to the edge device which has become disconnected from the first trunk switch.

13 Claims, 5 Drawing Sheets

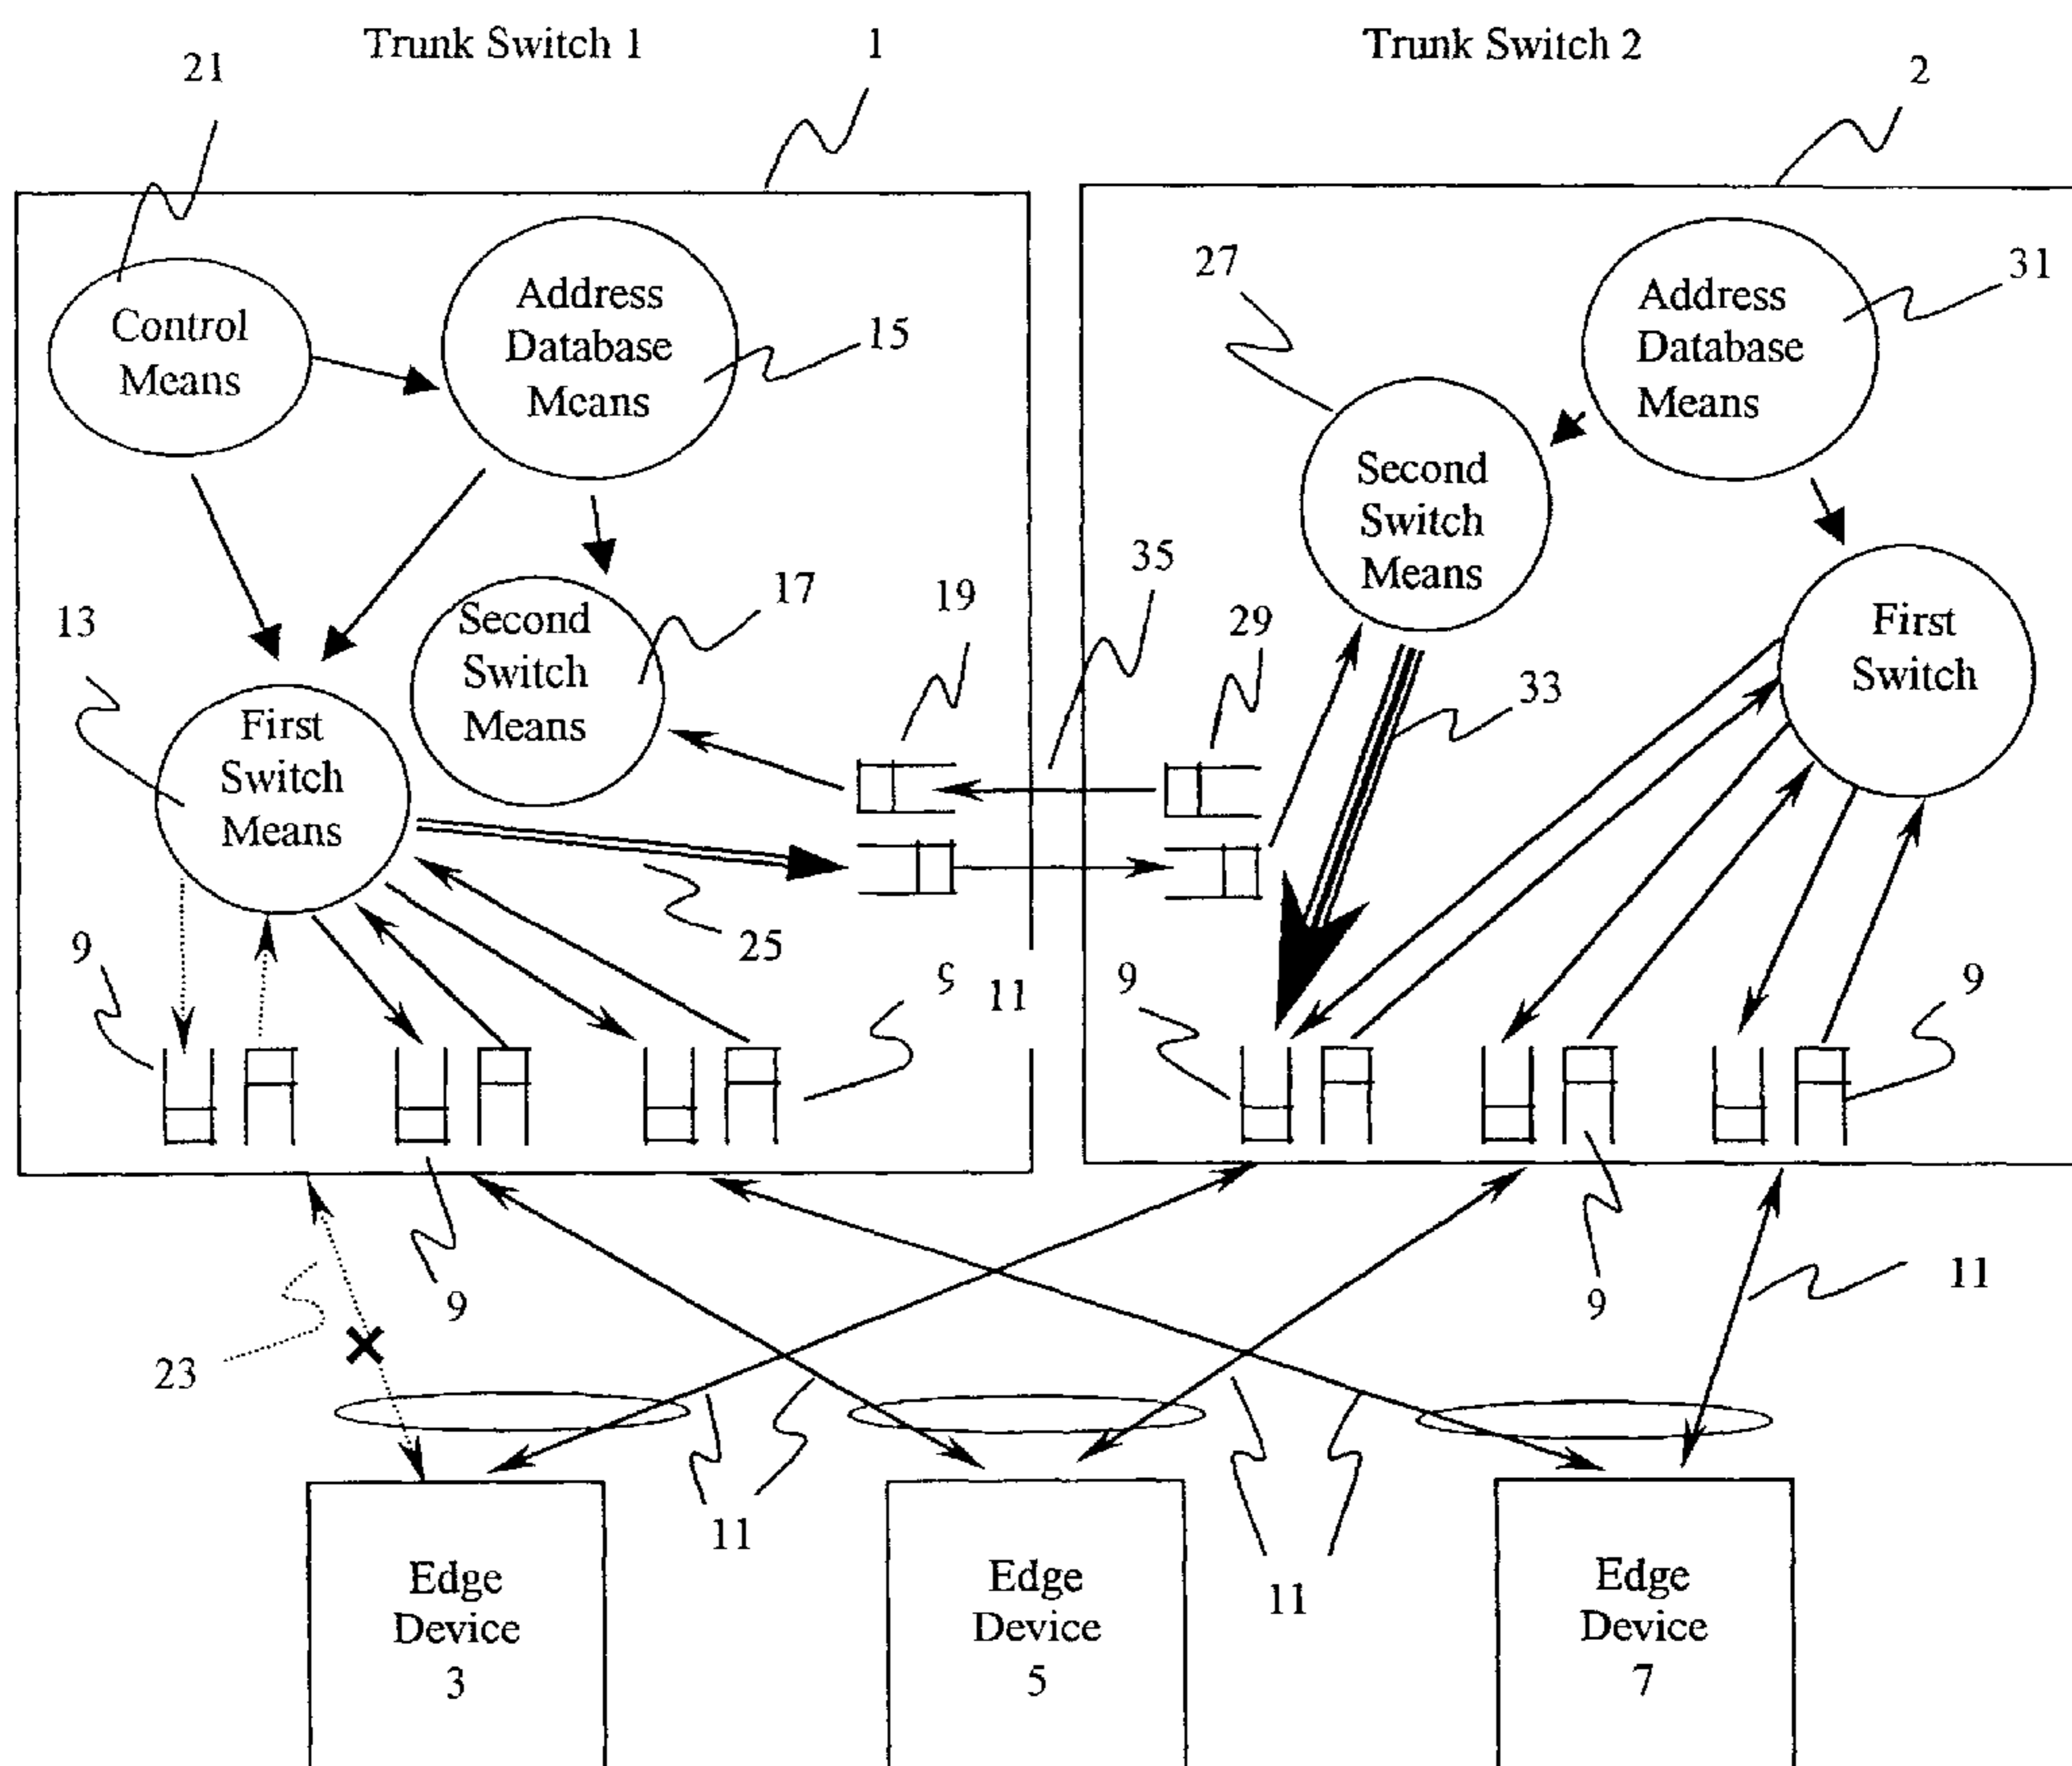


FIGURE 1

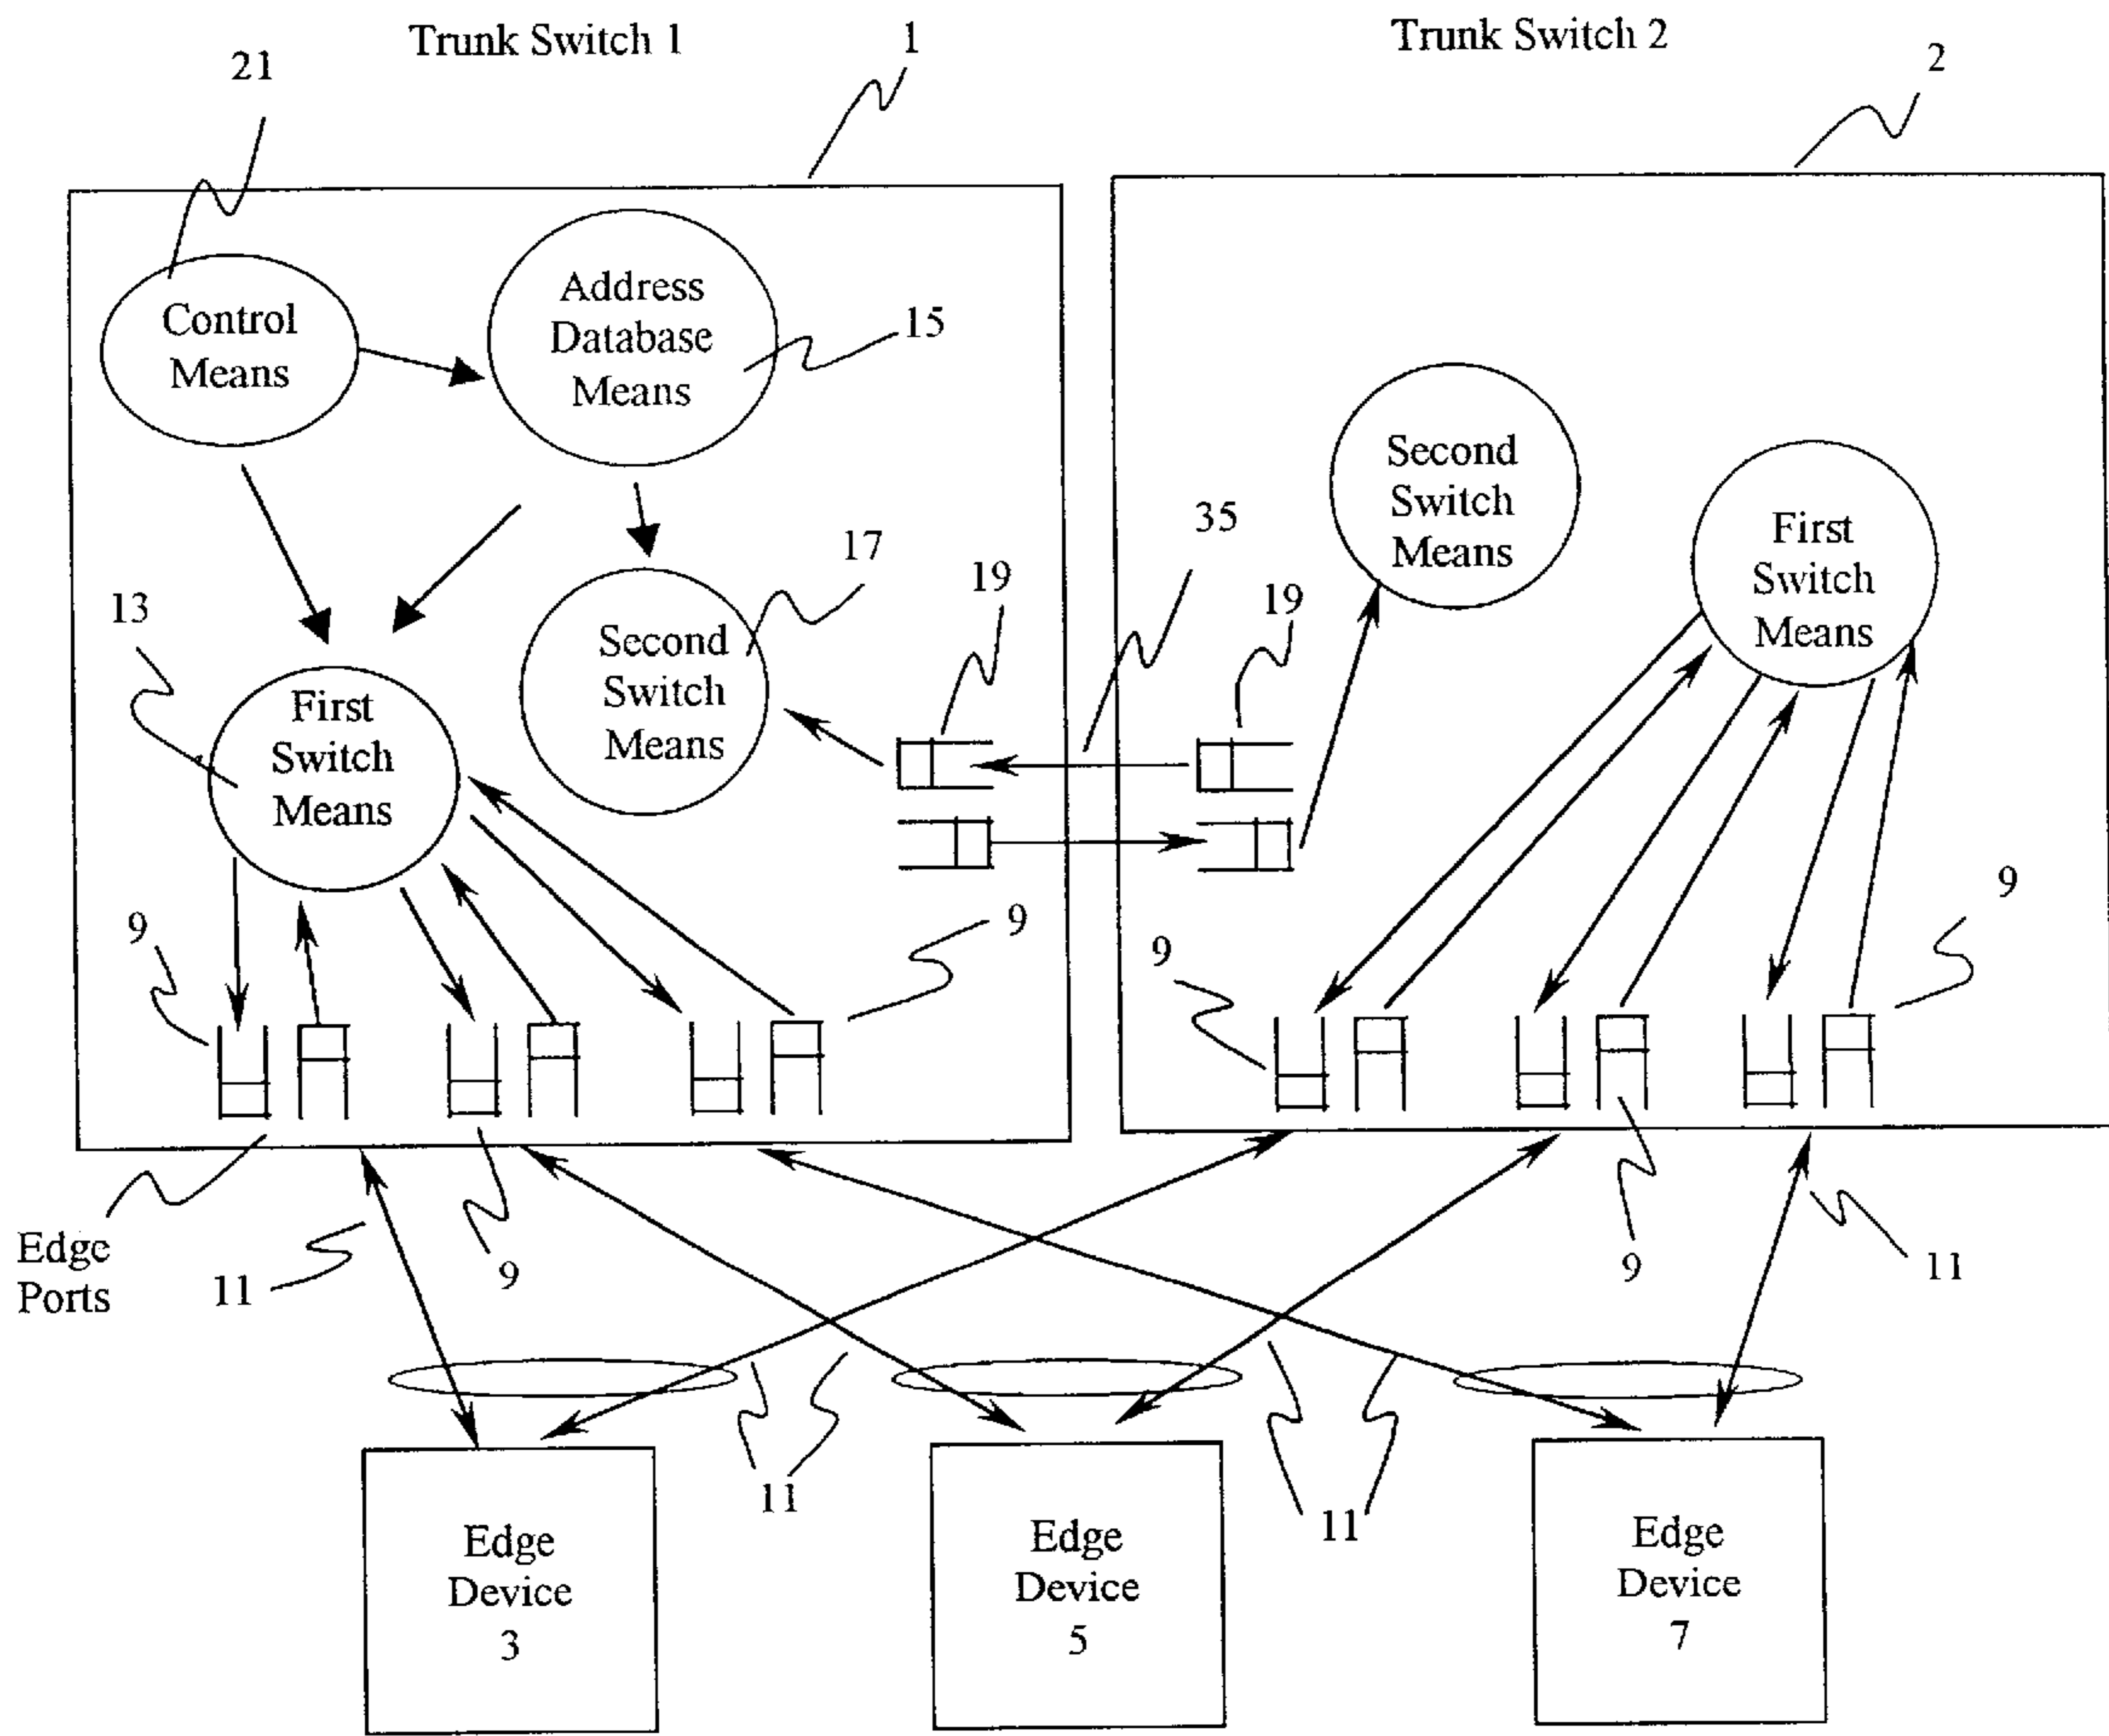


FIGURE 2

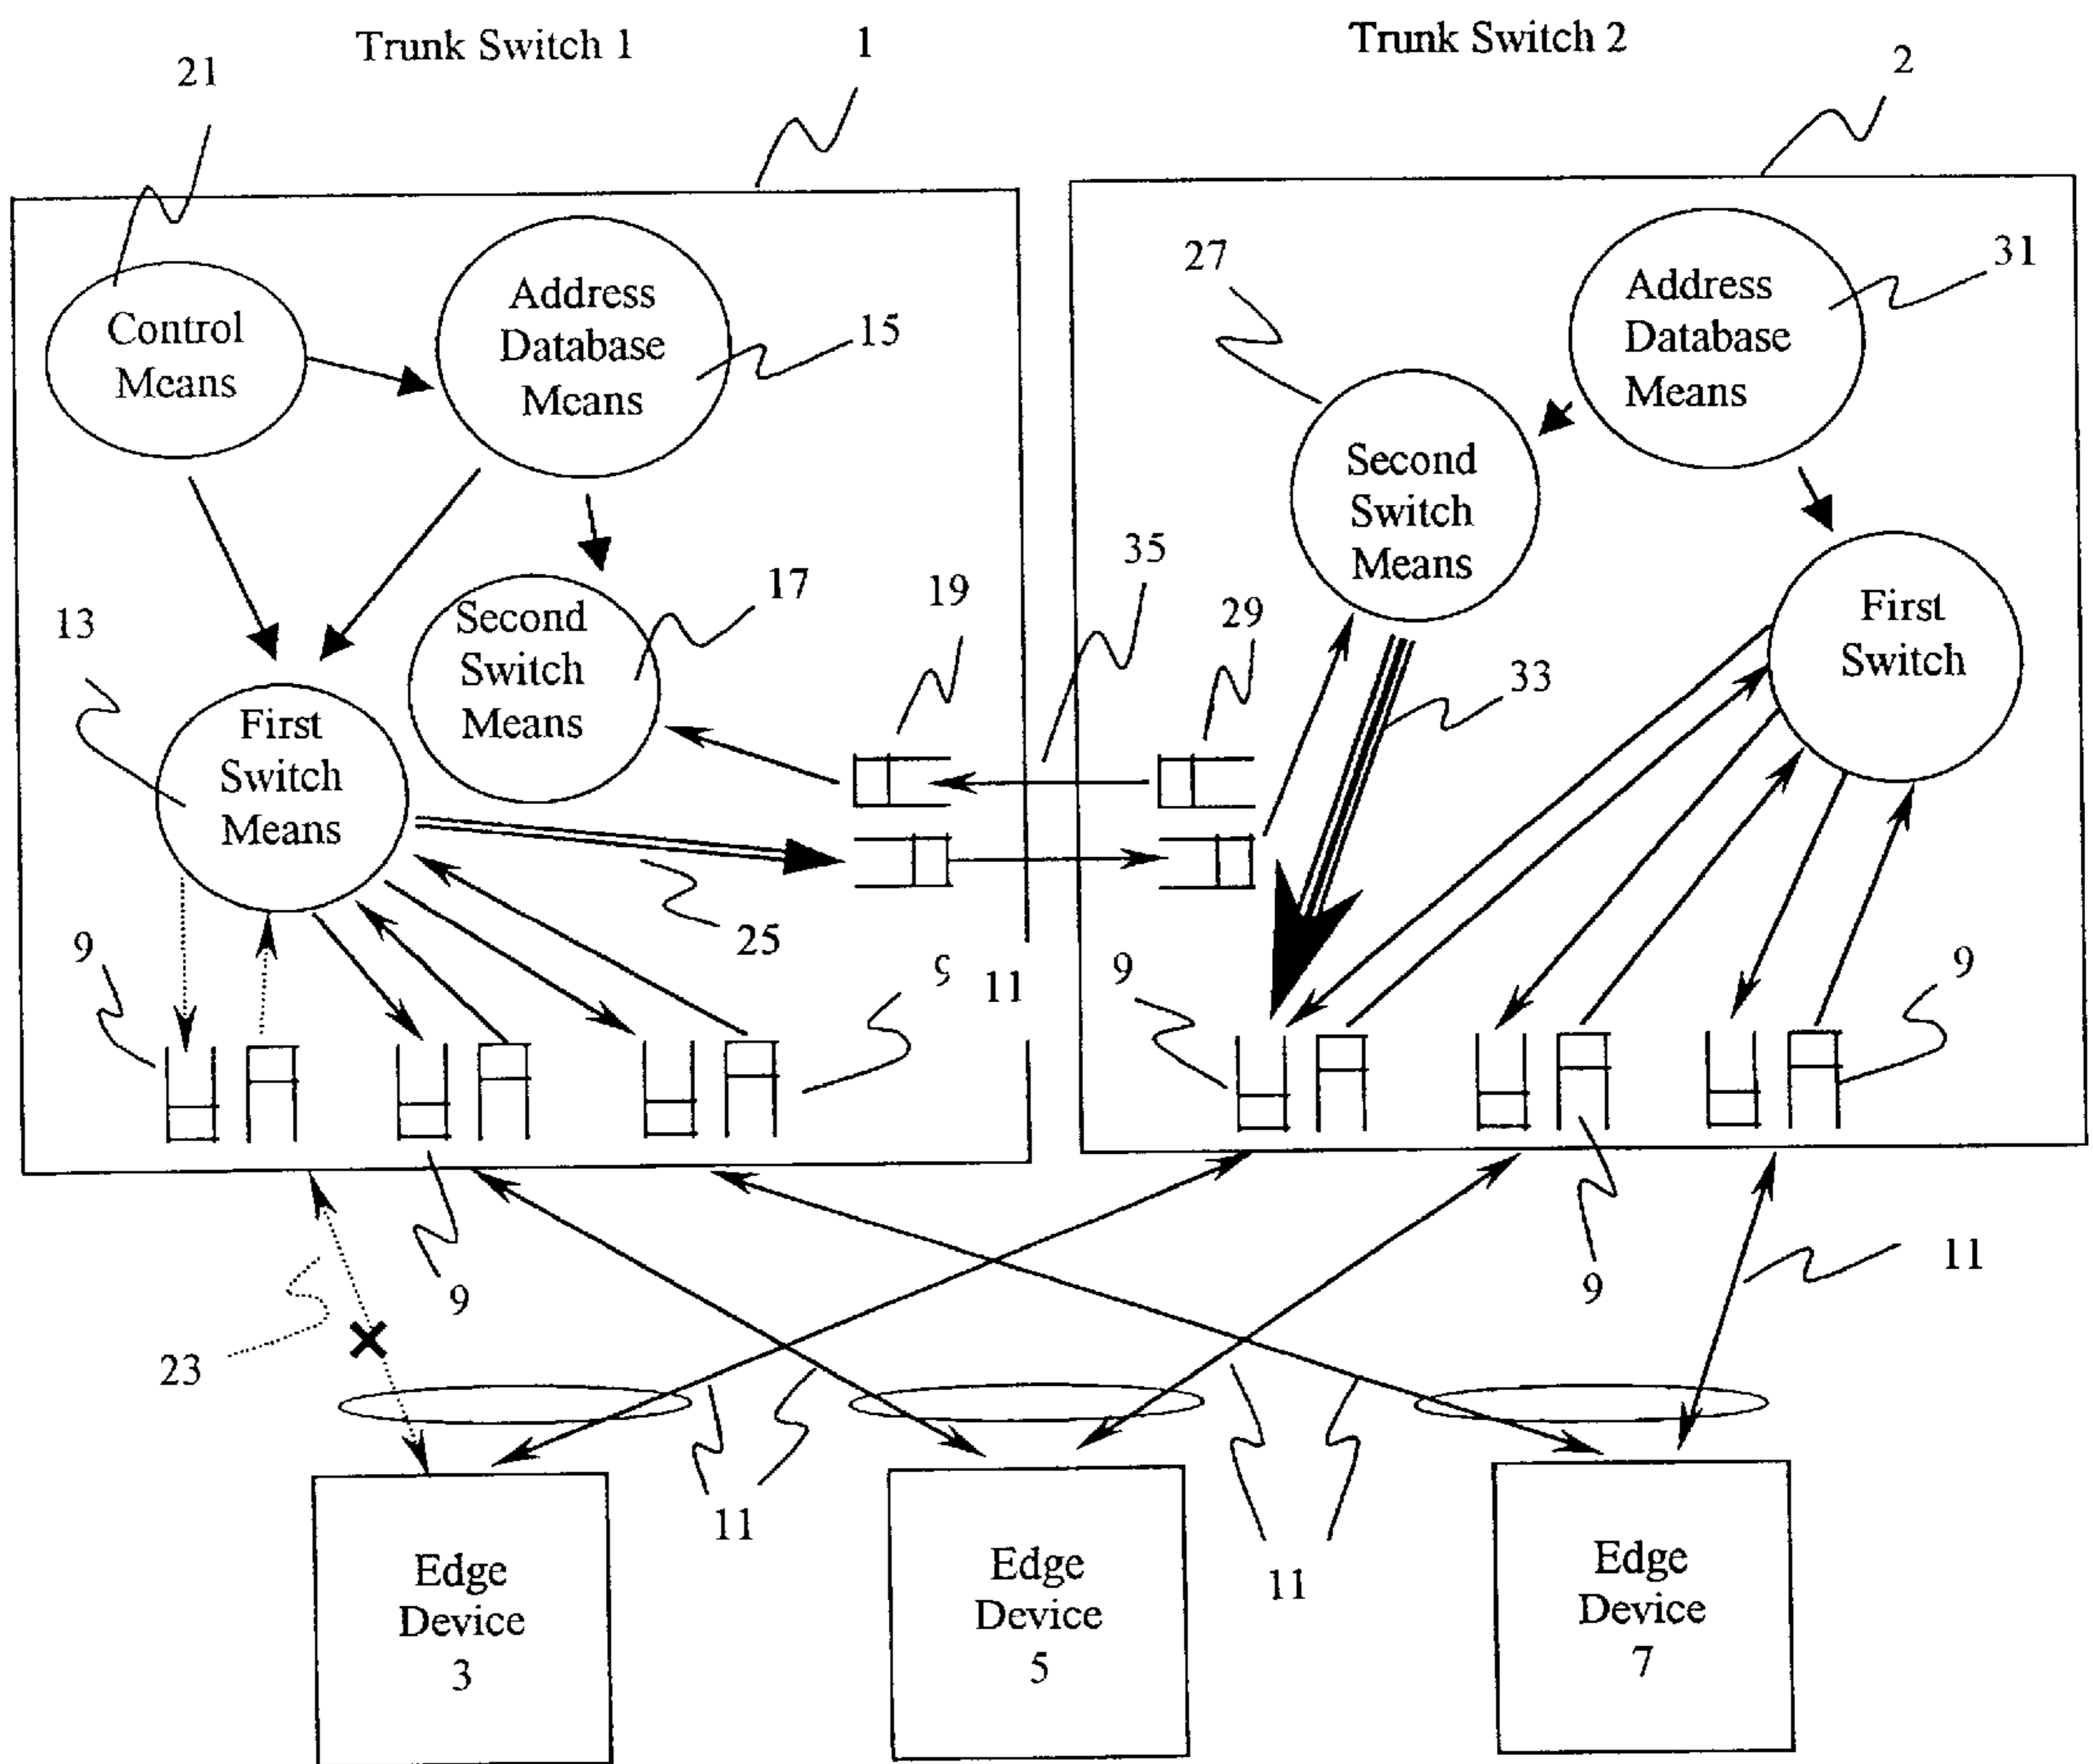


FIGURE 3

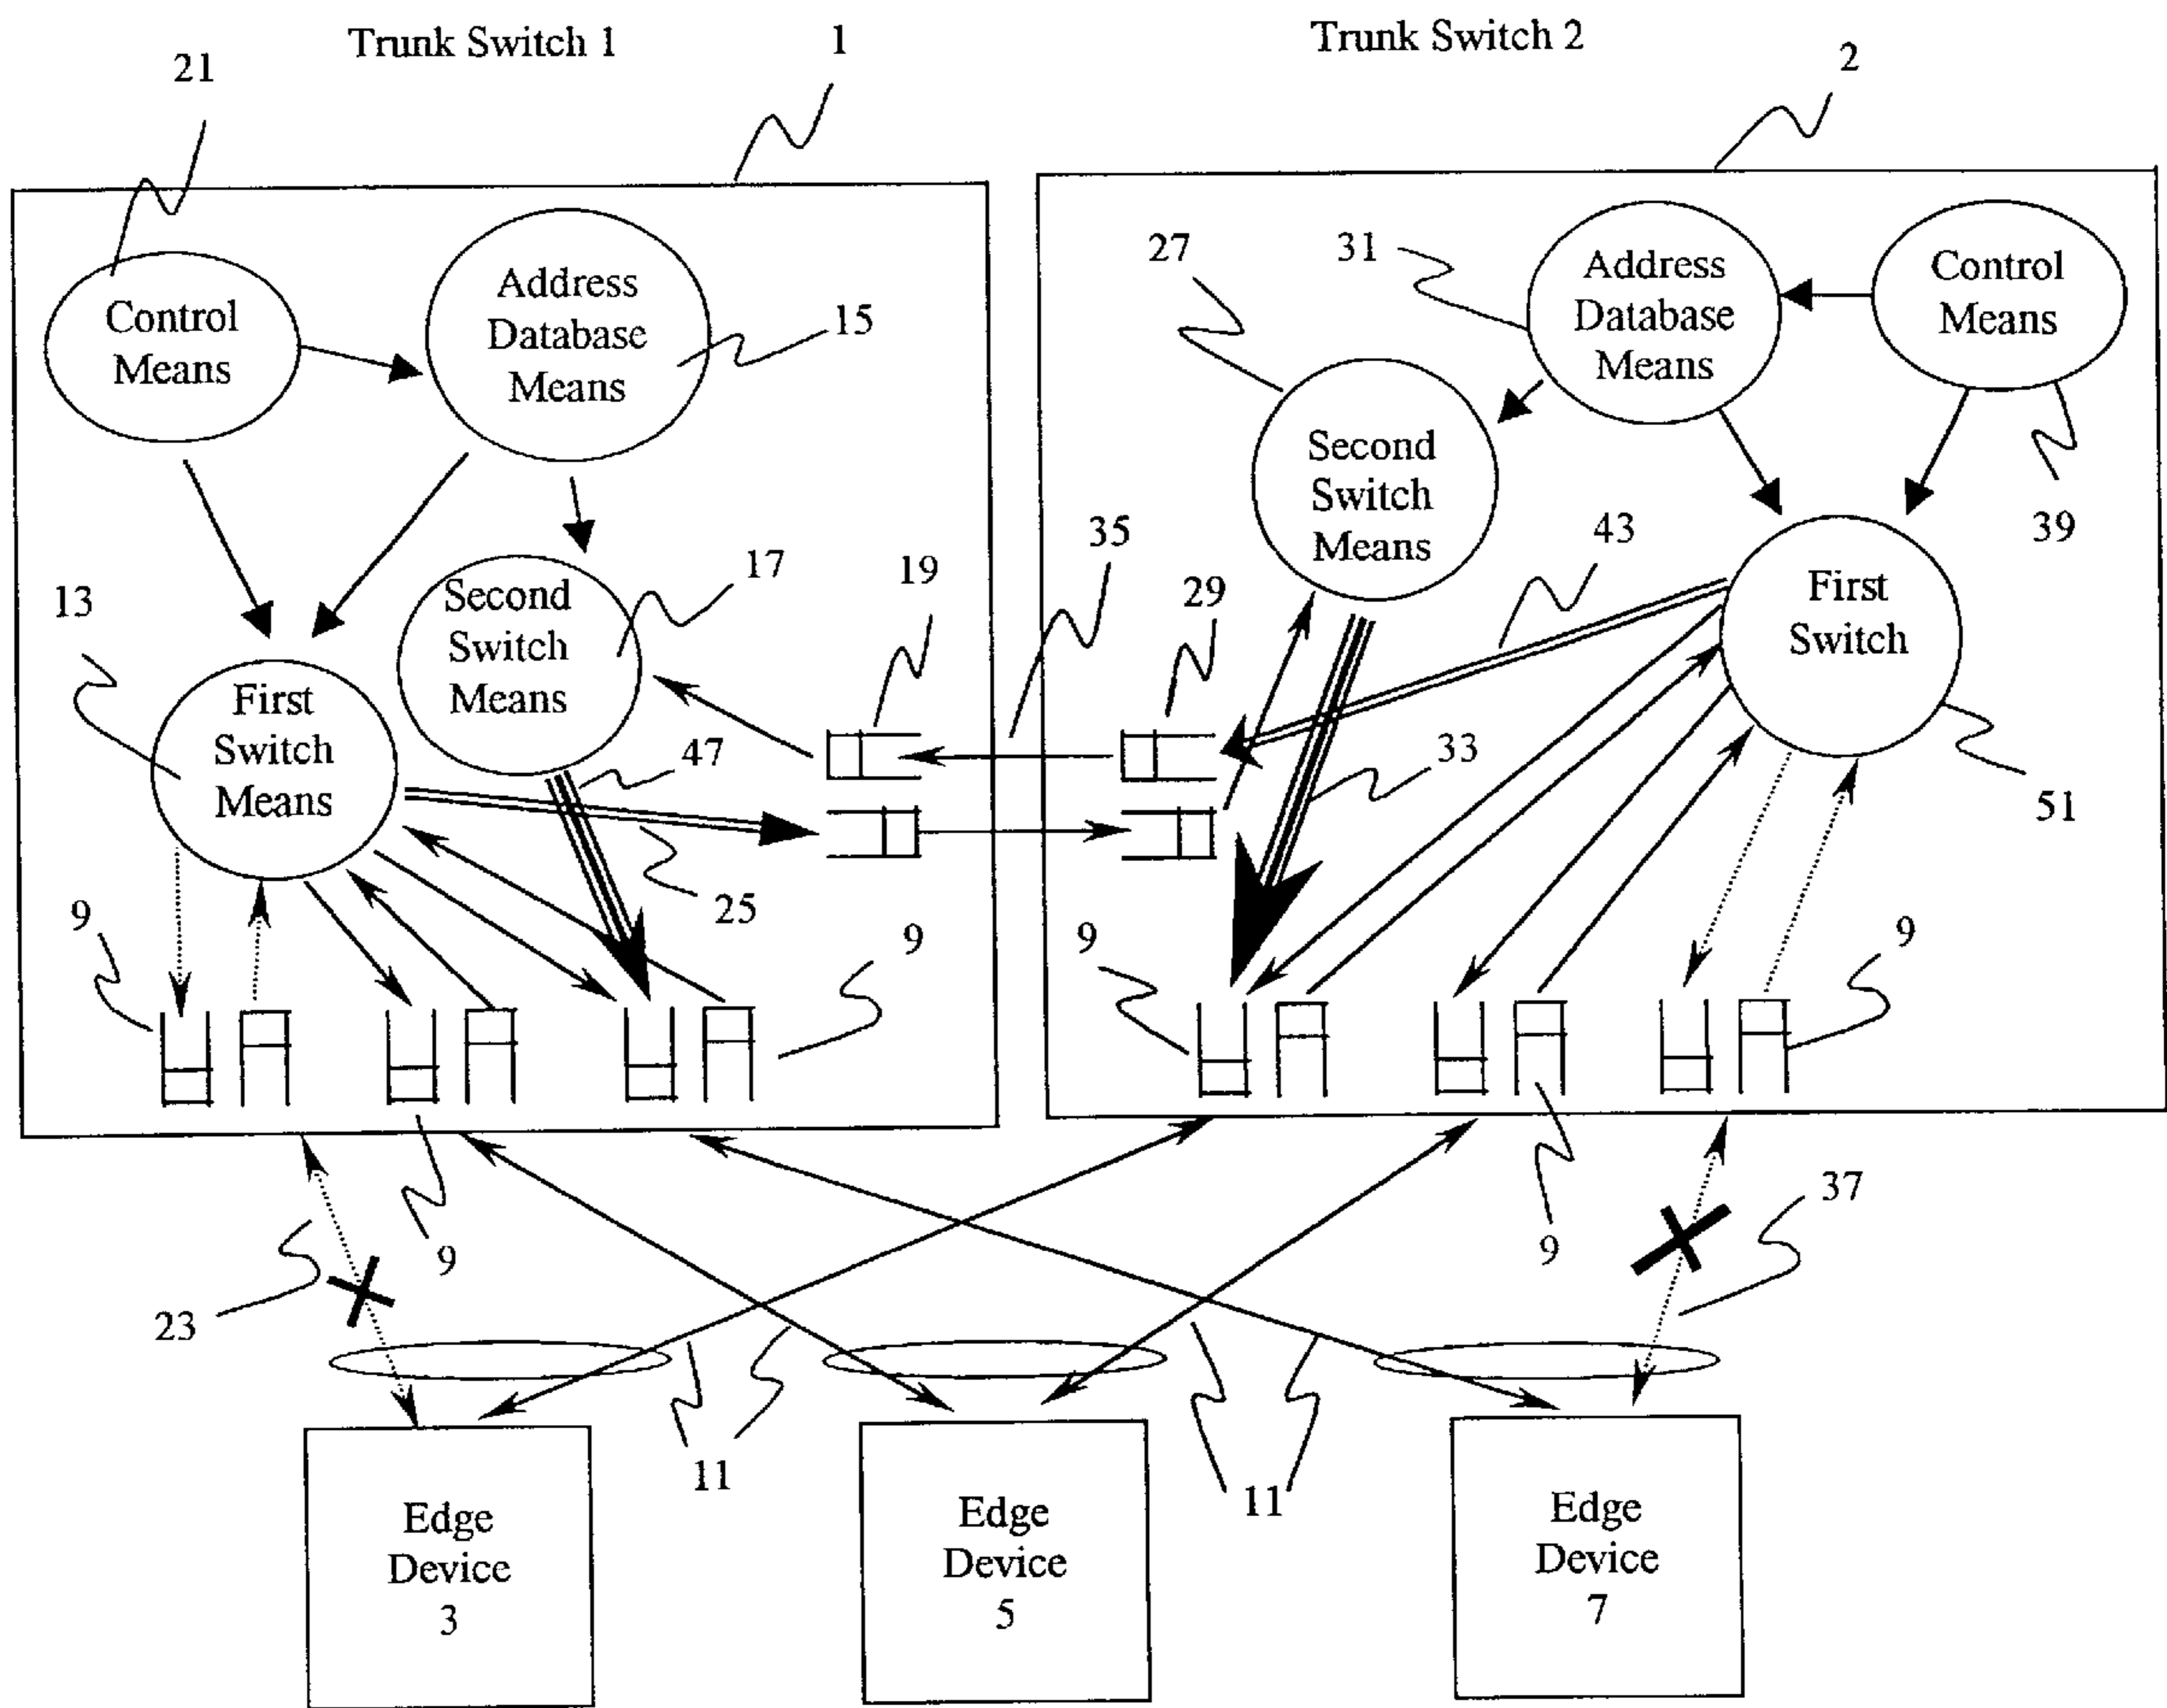


FIGURE 4

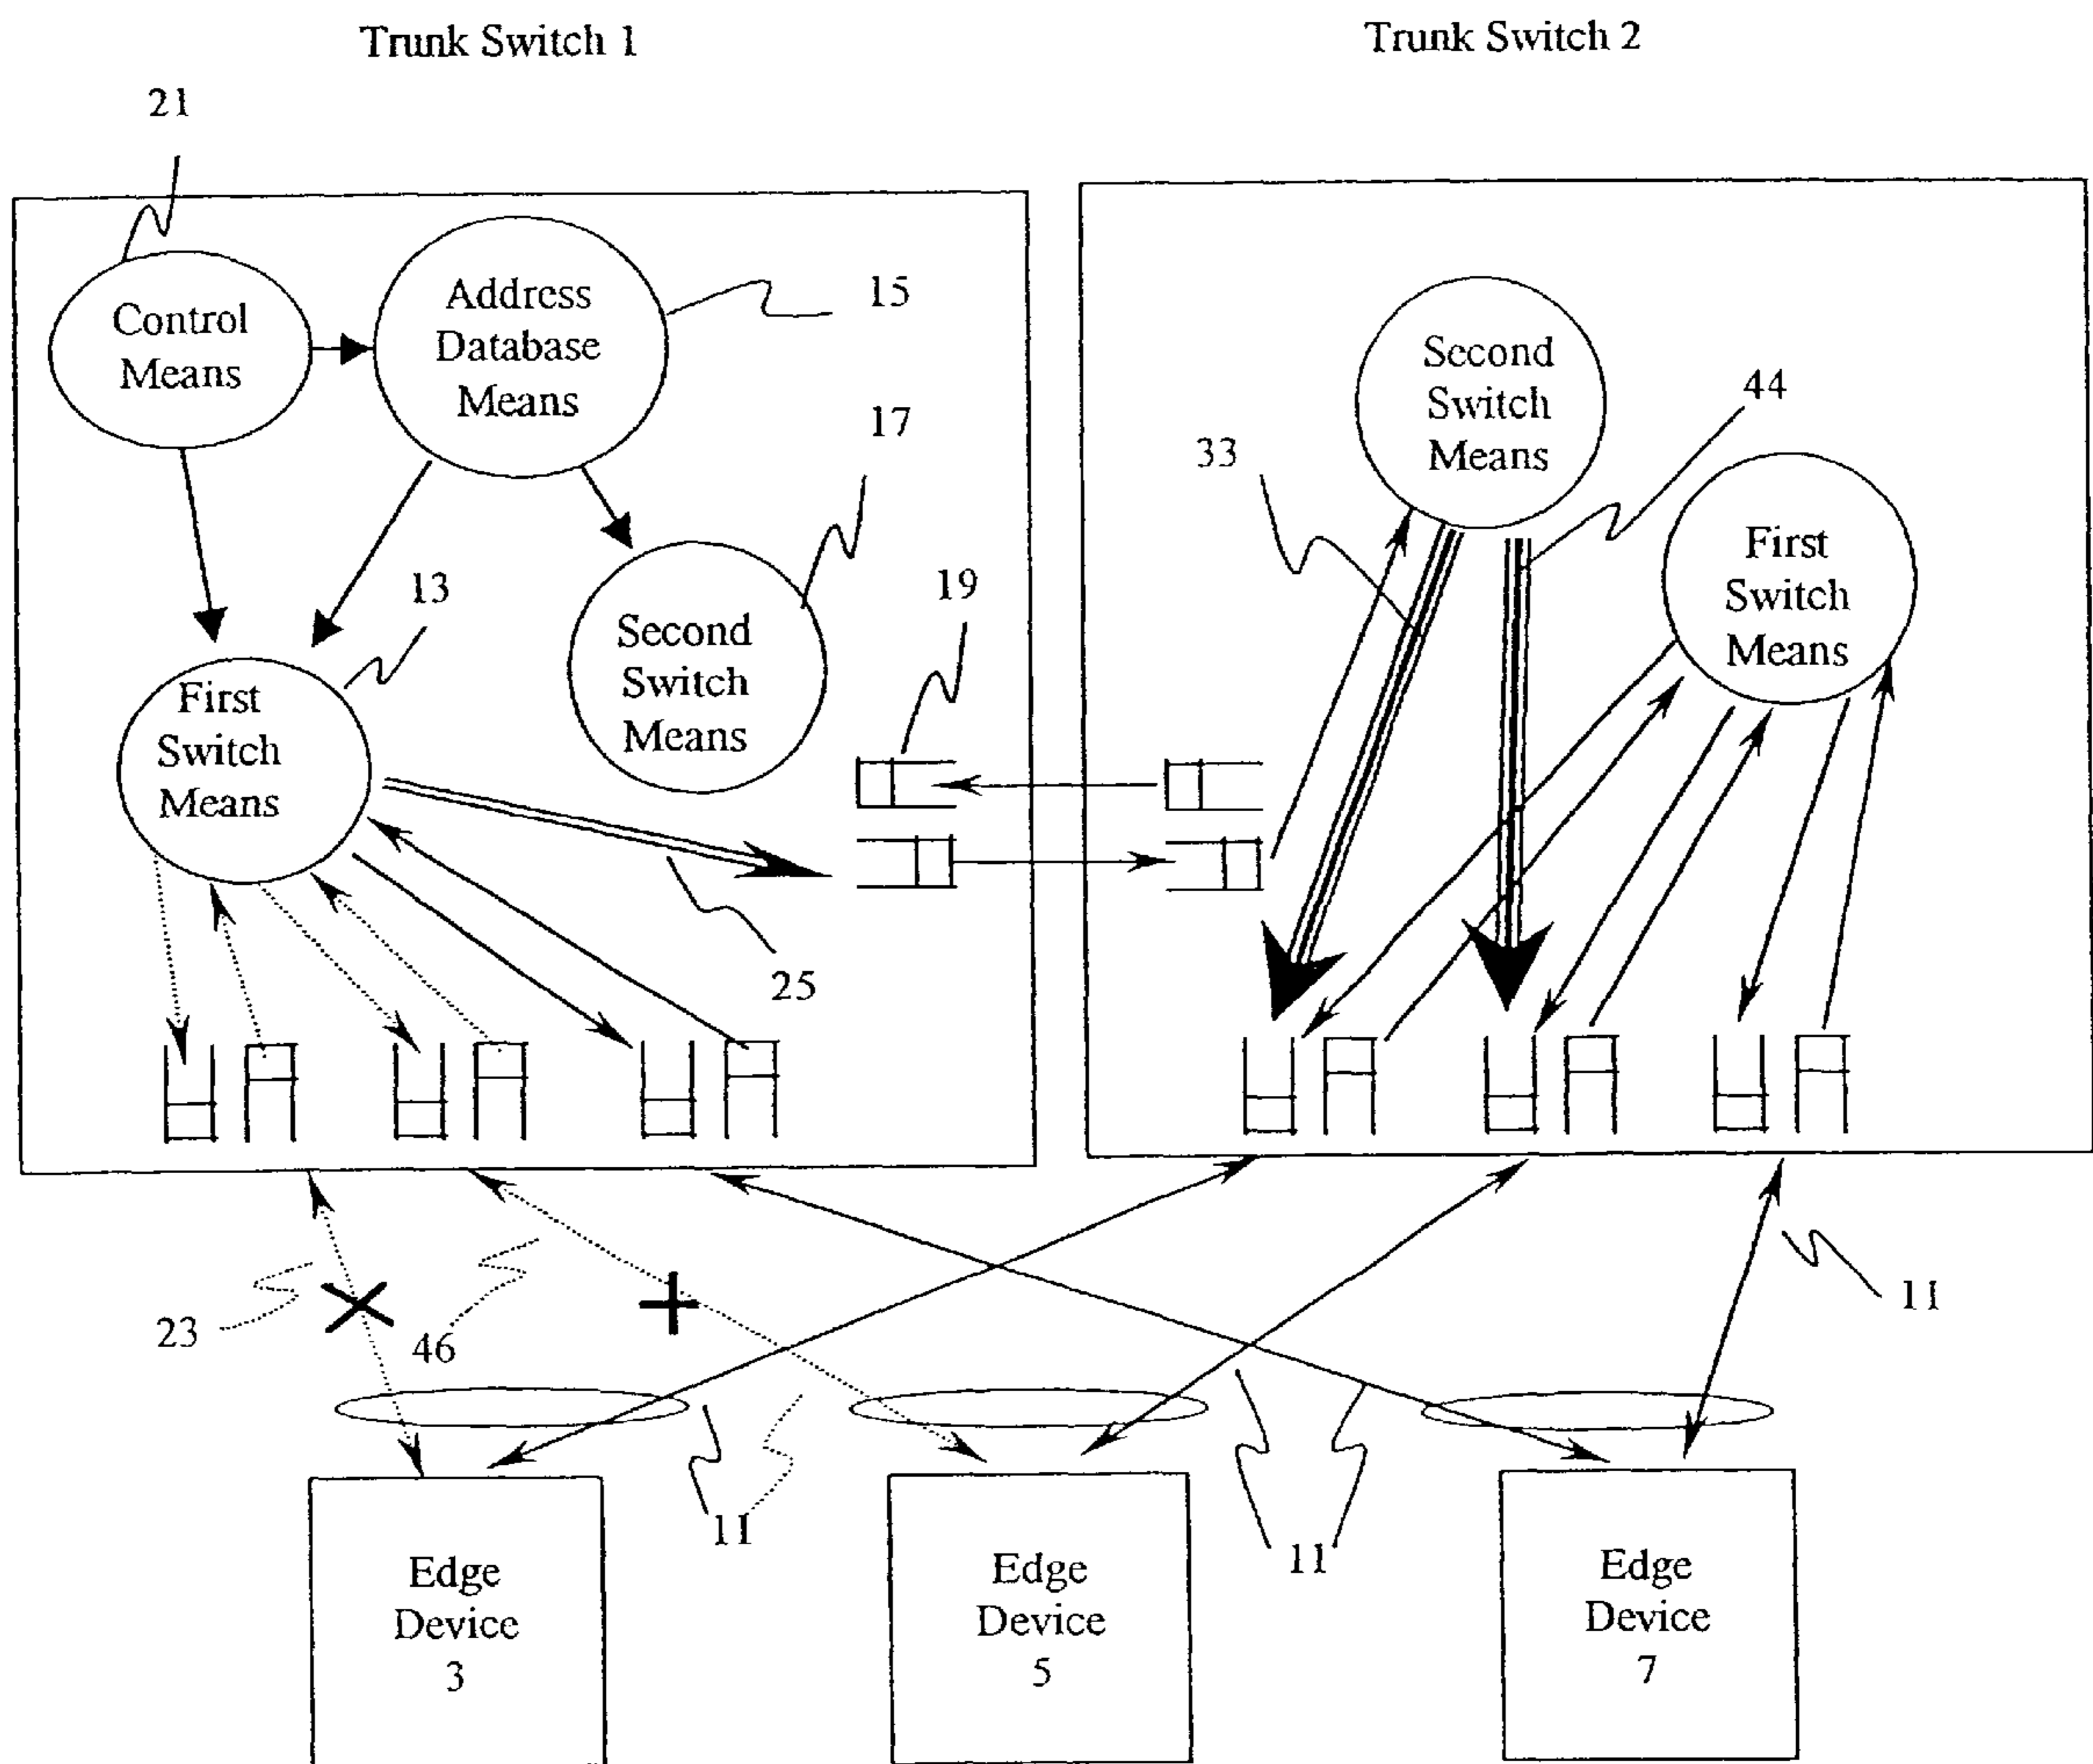
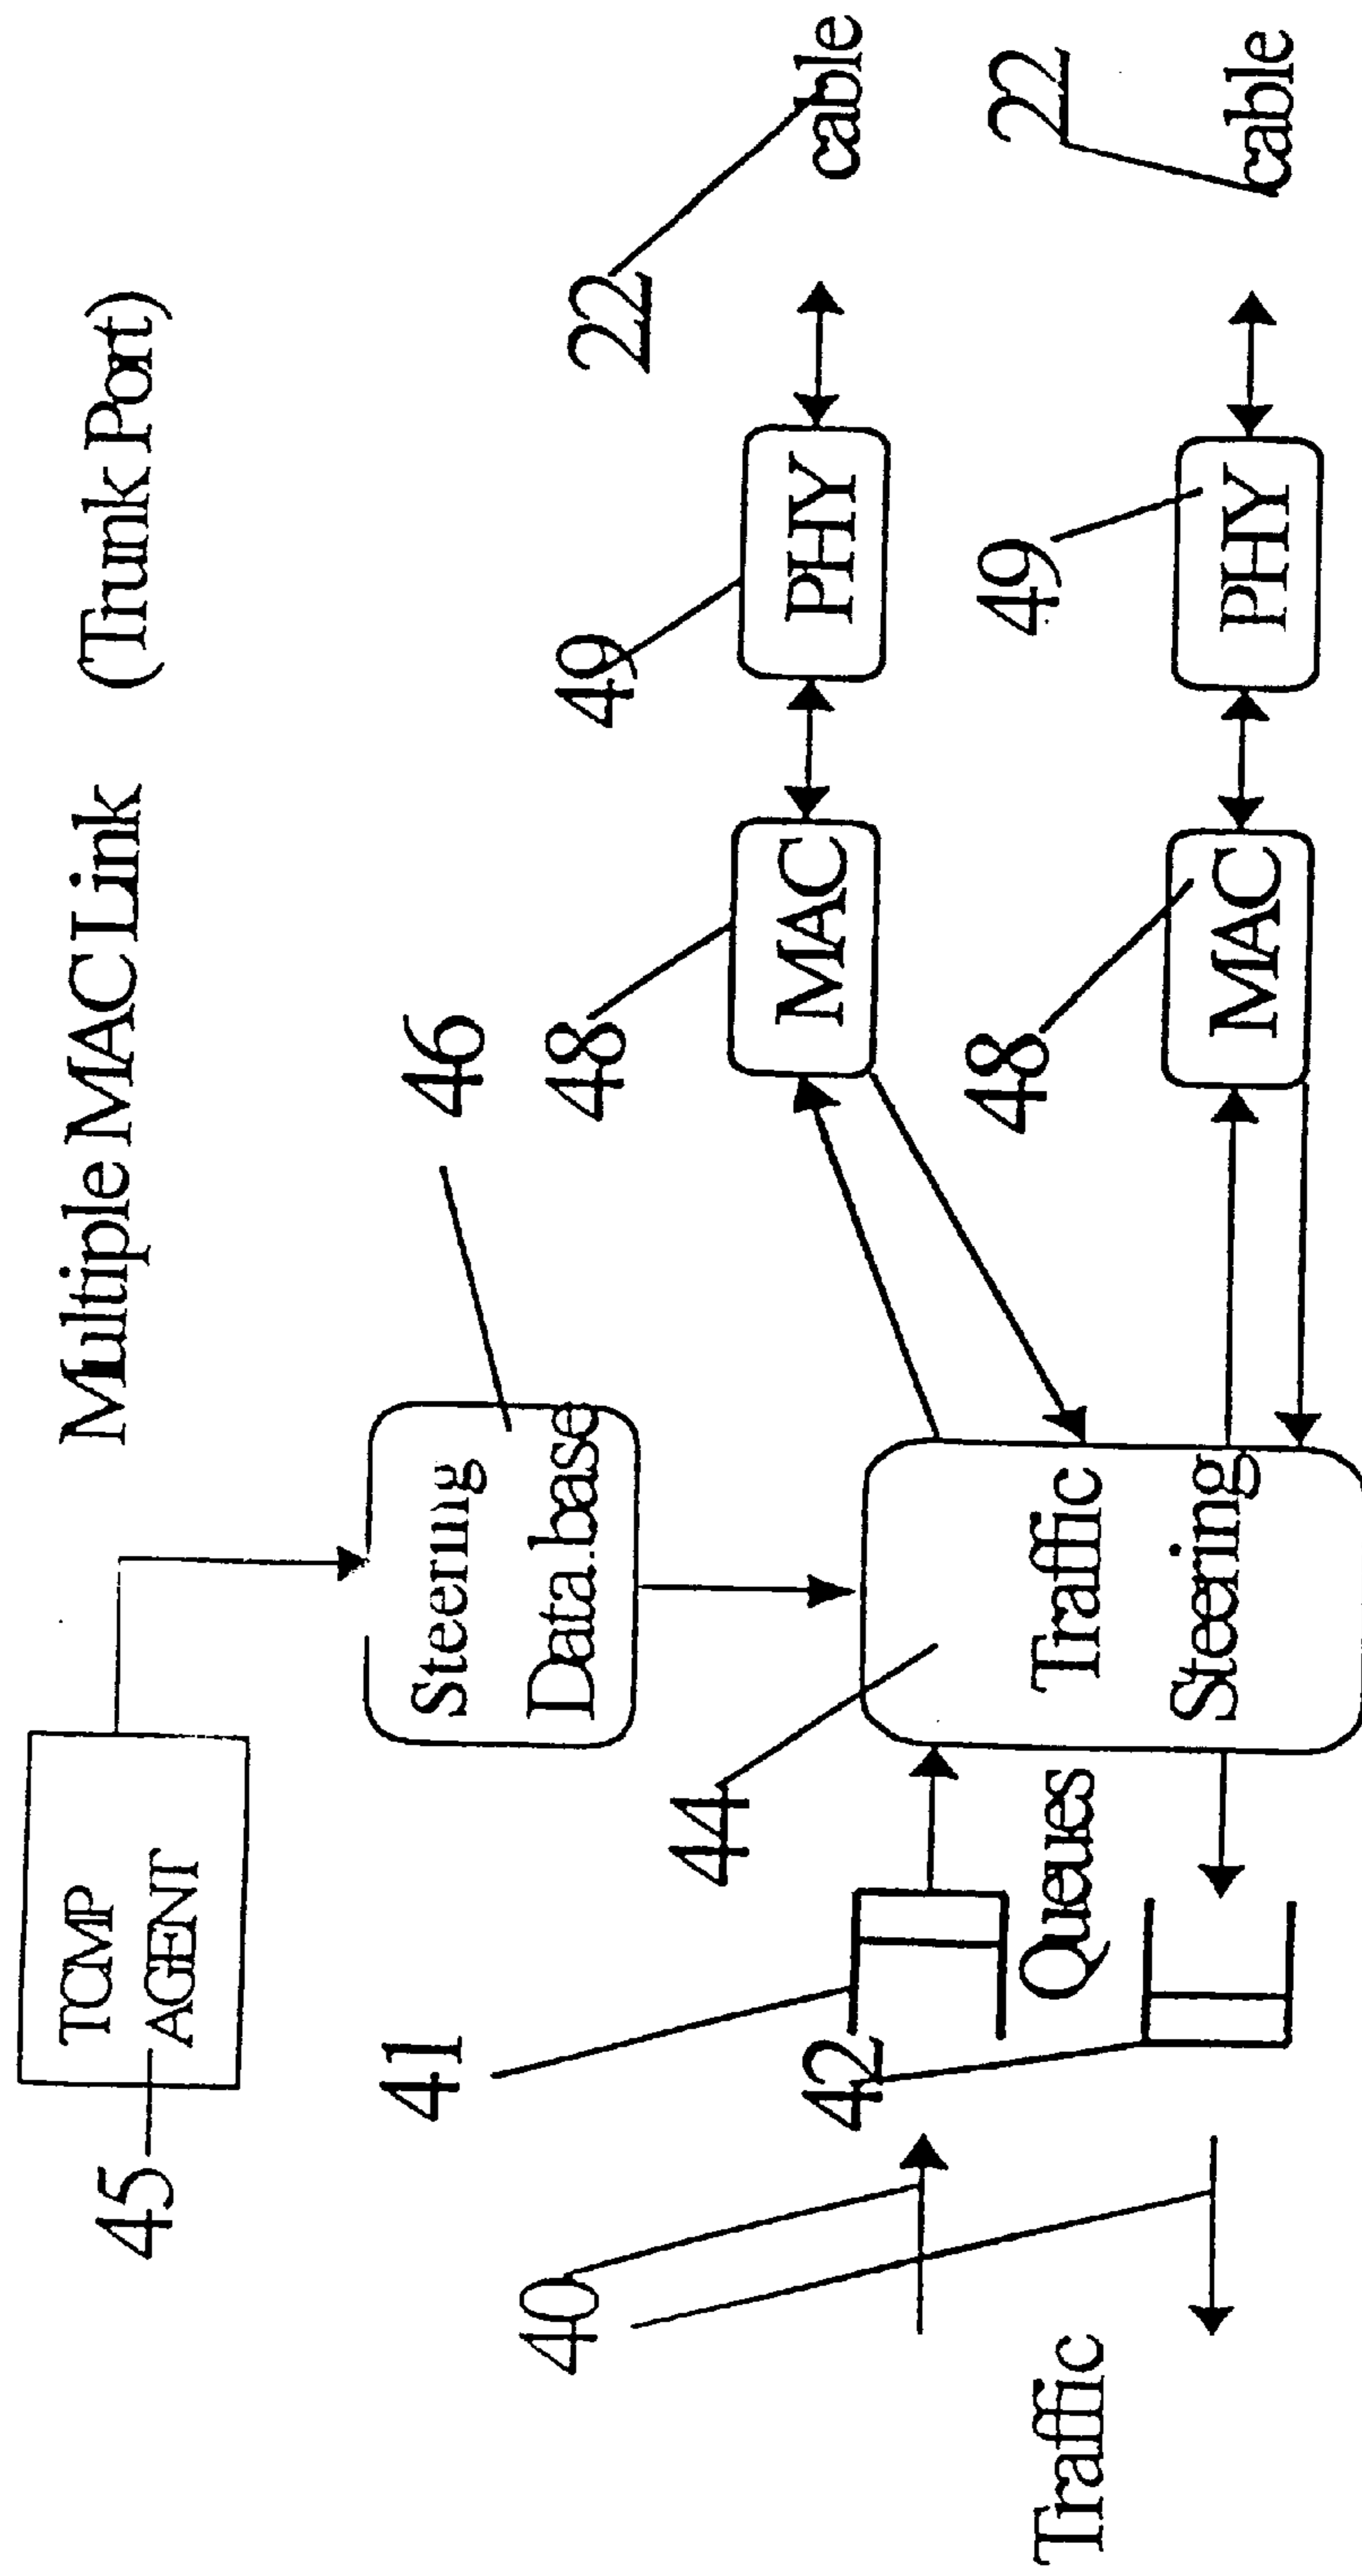


FIGURE 5



TRUNK ARRANGEMENT IN A COMPUTER NETWORK WITH INTRA TRUNK SWITCH CONNECTIONS

FIELD OF THE INVENTION

The present invention relates to connecting a plurality of edge devices to each other over a plurality of lines, and in particular through a plurality of trunk switches where the trunk switches have a high speed connection between them.

BACKGROUND OF THE INVENTION

In a computer network, devices communicate with each other over links. The amount of information or data that can be transmitted over a link, is limited by the bandwidth of the link. Different types of links, such as twisted pair, shielded twisted pair, coaxial and fiber-optic, have different amounts of bandwidth. Generally the larger the bandwidth of a link, the more expensive the cost per length of the link, and the difficulty in using the link. Presently, the rate at which information or data is desired between devices is much higher than the bandwidth of existing links, or economically practical links. Furthermore, it is foreseen that the desired data rates will constantly and significantly increase, while increasing the bandwidth of a link is seen to be much more difficult in the future.

One possibility for increasing the data rate between two devices is to use a plurality of links. This is especially beneficial in computer networks where the information to be transported is divided into a plurality of packets. The packets are then divided over the plurality of links and the data rate of the connection is substantially the sum of the data rate that each link can transfer.

The plurality of combined links are known as trunked lines or links. Trunked lines or links are useful when two separate devices need to communicate directly with each other at high data rates.

However, if three or more devices, and possibly tens or hundreds of devices, desire, or require, a high data rate connection, the use of a plurality of trunk lines by themselves would not be particularly effective.

Applicant's co-pending patent applications Ser. Nos. 9/014,548, 09/014,547, 09/060,606 and 09/114,975, and are incorporated by reference. In these applications, a trunking arrangement is described where each device is connected to a plurality of trunk switches. The devices are called edge devices and the data to be transferred is divided into packets with destination addresses.

Each trunk switch has a plurality of ports for connecting to all of the edge devices. Packets received on a port are analyzed to determine which port the packet should be sent out on to reach its destination device. Each trunk switch can therefore receive packets from all the edge devices and properly forward the packets to all the edge devices.

An edge device can therefore divide its packets of information among the trunk links and switches. All the packets will then reach the proper destination device, and the data rate will be the sum of the data rate of all the trunk links and switches.

The links connecting each edge device to each trunk switch, and the plurality of trunk switches form a packet based high speed mesh, called a trunk cluster. As described in the co-pending patent applications, this trunk cluster can be constructed with a set of loosely coupled switches, a configuration protocol, trunk network interfaces and optionally a reachability protocol.

Each switch, called a Trunk Switch, in a trunk cluster, provides a single "shared LAN" by interconnecting two or more links. The attached edge devices connect to a Trunk Cluster via trunked links with either basic trunking or trunking enhanced with station reachability. Basic Trunking requires at least one physical connection to each switch (1 through N) in the Trunk Cluster. As described in application Ser. No. 09/014,548 each switch in the trunk cluster provides a single shared LAN by interconnecting two or more links. The edge devices attached to the links run a trunk configuration protocol. These attached edge devices view each physical link within a trunked port as if the attached device is connected to a shared logical LAN with multiple other attached devices. Each Trunk Switch maintains a list of identifiers that uniquely identify each edge device that is connected to it. The identifier is called the Trunk Port Identifier (TPI). It builds this list up by monitoring the trunk configuration protocol packets that are transmitted by the edge devices. The trunk switch builds TPI to Link map that relates the switch's link identifiers to the TPIs.

A difficulty with these trunk cluster occurs when one of the edge devices loses its connection to one of the trunk switches. In order to maintain the trunking arrangement operational for all edge devices, the trunk switch having the connection problem is effectively removed from the trunk cluster. This allows the trunking arrangement to continue proper operation. However losing this trunk switch causes the bandwidth of the trunk cluster to decrease, which correspondingly lowers the data rate of all the edge devices, due to one edge device losing a link to one trunk switch.

SUMMARY AND OBJECTS OF THE INVENTION

It is a primary object of the present invention to provide a trunking arrangement with a plurality of trunk switches, where a failure of one or more links from an edge device to a trunk switch does not require that the entire trunk switch be removed from the trunking arrangement. The present invention allows the trunk switch to remain in the trunk cluster and to continue forwarding packets of data.

The present invention accomplishes this objective by providing a particular type of trunk switch in the trunk cluster. This trunk switch has a plurality of edge ports for separately connecting to each of the plurality of edge devices by a separate edge link. An address database means is located in the trunk switch for indicating a correspondence of the edge ports to the edge devices. A first or normal switch means in the trunk switch reads the data from each of the edge ports, and consults the address database means to determine which edge port to send the received data out on. The first switch means then transfers the received data to the determined edge port.

The trunk switch of the present invention also includes an intra-cluster port for connecting the individual trunk switches to each other through an intra-cluster link. A second or uni-direct switch means reads intra-cluster data from said intra-cluster port and consults the address database means to determine which edge port to send the intra cluster data out on. The second switch means then transfers the intra-cluster data to the determined edge port. The trunk switch also has a control means for detecting when an edge device has become disconnected. The control means then modifies the address database means to redirect data for the disconnected edge device to the intra-cluster port for transmission to another trunk switch.

The edge ports and the intra-cluster port transfer the data packets in a substantially identical manner. Also the data

packets transferred over the intra-cluster ports and links are transferred at substantially the same data rate as the data packets transferred over the edge ports and edge links.

In particular the first switch means creates a transmit only connection from the first switch means to the intra-cluster port when the control means detects said disconnected edge device. Correspondingly the another trunk switch creates a transmit only connection from the second switch means to the port of an edge device which corresponds to a disconnected edge device of another trunk switch which sent the intra-cluster data.

Since each trunk switch can forward data packets for a disconnected edge device to another trunk switch, and an edge device which is disconnected from one trunk switch can still transmit to another trunk switch, the trunk switch having a disconnected edge device does not need to be removed from the trunking arrangement of the present invention. Therefore while the data rate will be reduced for the edge device which has been disconnected from the trunk switch, the data rates for the remaining edge devices will not be significantly reduced, if reduced at all. The present invention is therefore able to provide a trunking arrangement which is more beneficial in the case of a failure of one of the edge links, and is an improvement in computer networking.

The various features of novelty which characterize the invention are pointed out with particularity in the claims annexed to and forming a part of this disclosure. For a better understanding of the invention, its operating advantages and specific objects attained by its uses, reference is made to the accompanying drawings and descriptive matter in which a preferred embodiment of the invention is illustrated.

BRIEF DESCRIPTION OF THE DRAWINGS

In the drawings:

FIG. 1 is a diagram of a trunking arrangement with all the edge devices connected to all of the trunk switches;

FIG. 2 is a diagram of a trunking arrangement with one of the edge devices disconnected from one of the trunk switches;

FIG. 3 is a diagram of a trunking arrangement where one of the edge devices is disconnected from one of the trunk switches and another edge device is disconnected from another of the trunk switches;

FIG. 4 is a diagram of a trunking arrangement with two of the edge devices disconnected from one of the trunk switches; and

FIG. 5 is a diagram showing traffic flow through a multiple link trunked port.

DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring to the drawings, in particular to FIG. 1, the trunking arrangement consists of trunk switches 1 and 2, and edge devices 3, 5 and 7. Each trunk switch 1, 2 contains a plurality of edge ports 9 which connect to the edge devices 3, 5 and 7 by edge links 11. Edge ports 9 operate according to IEEE 802.1d standard for MAC clients.

Each trunk switch 1, 2 includes a first switch means which transfers data packets received on one edge port 9 to another edge port 9. The first switch means 13 consults an address database means 15 for determining which other edge port 9 the received data should be sent out on. The second switch means 17 reads data from an intra-cluster port 19. A control means 21 monitors the first switch means 13 and address database means 15. The control means 21 is also able to modify the first switch means 13 and the address database means 15.

An edge device 3 can become disconnected from one of the trunk switches 1, 2 due to many different reasons. As shown in FIG. 2, edge device 3 can become disconnected from trunk switch 1 due to a defective edge link 23. The control means 21 detects from the first switch means 13 that edge device 3 is no longer communicating properly with trunk switch 1. Control means 21 then modifies the address database means 15 so that the first or normal switch means 13 will send packets destined for edge device 3 to intra-cluster port 19 instead of the port connected to the defective edge link 23. The first switch means 13 then creates a transmit only connection (double line arrow) 25 to the intra-cluster port 19. The first trunk switch 1 also communicates to the trunk switch 2 that the link 23 from trunk switch 1 to edge device 3 is non-operational. The second switch means 27 of trunk switch 2 will therefore create a transmit only connection (triple line arrow) 33 from the second switch means 27 to the edge port 9 which connects to the edge device 3.

Second switch means 27 of trunk switch 2 monitors the intra-cluster port 29 of trunk switch 2 and receives the data packets that first switch means 13 sends to intra-cluster port 19 since edge device 3 was disconnected from trunk switch 1. Second switch means 27 consults the address database means 31 to determine which edge port of trunk switch 2 the intra-cluster data-packet should be sent to. Since trunk switch 1 and trunk switch 2 operate according to identical algorithms, and their address databases means 15, 31 are synchronized, both trunk switches will determine that this data-packet is destined for edge device 3. In this way, when trunk switch 1 determines that it cannot deliver a packet, the packet is then sent to trunk switch 2 through the intra-cluster port 19 and intra-cluster link 35 to trunk switch 2. Trunk switch 2 then processes this intra-cluster data packet and forwards the intra-cluster data packet to the proper edge port according to the standard processing of data packets and the transmit only connection 33 which is created when another trunk switch loses a connection.

The operations of the trunk switches is described by the following pseudo code:

```
If edge device 3 loses its connection to trunk switch 1 and
maintains its connection to trunk switch 2 then
  Trunk switch 1 adds a transmit only connection 25
  (double line arrow) from the first switch means 13 to
  the intra-cluster port 19;
  Trunk switch 1 modifies those entries in address data
  base means 15 which direct packets to the port
  connected to edge device 3 to direct packets to the
  intra-cluster port 19;
  Trunk switch 2 adds a transmit only connection 33
  (triple line arrow) from the second switch means 17
  to the edge port 9 connected to edge device 3.
```

The trunk cluster of the present invention is also able to adapt to a situation where two separate edge devices 3 and 7 become disconnected from two separate trunk switches 1 and 2. Trunk switch 1 detects the defective edge link 23 as previously described and the first switch means 13 forwards the data for the edge device 3 to intra-cluster port 19 instead of the edge port 9 which corresponds to edge device 3. Trunk switch 2 receives the intra-cluster data over intra-cluster port 29 and the second switch means 27 transfers this intra-cluster data to edge port 9 which corresponds to edge device 3.

Control means 39 of trunk switch 2 detects defective link 37 in the same manner as control means 21 of trunk switch 1 detects defective link 23. First switch means 51 forward data destined for edge device 7 to intra-cluster port 29 where

it is transmitted over intra-cluster link 35 to intra-cluster port 19. The second switch means 17 transfers the intra-cluster data from intra-cluster port 19 to the edge port 9 which corresponds to edge device 7. This is shown in FIG. 3.

In FIG. 3 edge device 3 loses its connection to trunk switch 1 and edge device 7 loses its connection to trunk switch 2. Each switch is running the pseudo code above, both compensate the lost links, and the changes for both errors work in together as described below.

When edge device 3 loses its connection 23 to trunk switch 1 and maintains its connection 11 to trunk switch 2, then; trunk switch 1 adds a transmit only connection (double line arrow) 25 from the first switch means 13 to the intra-cluster port 19. Trunk switch 1 modifies those entries in address data base means 15 which direct packets to the port connected to edge device 3 to direct packets to the intra-cluster port 19, and trunk switch 2 adds a transmit only connection (triple line arrow) 33 from the second switch means 27 to the port 9 connected to edge device 3.

In addition, when edge device 7 loses its connection 37 to trunk switch 2 and maintains its connection 11 to trunk switch 1, then; trunk switch 2 adds a transmit only connection (double line arrow) 43 from the first switch means 51 to the port 29. Trunk switch 2 modifies those entries in address data base means 15 which direct packets to the port connected to edge device 3 to direct packets to the intra-cluster port 29; and trunk switch 1 adds a transmit only connection (triple line arrow) 47 from the second switch 17 to the port 99 connected to edge device 7.

Using the pseudo code as previously described, trunk clusters of the present invention are also able to adapt when two or more edge devices have defective links to the same trunk switch. This situation is shown in FIG. 4 where edge device 3 has a defective link 23 to trunk switch 1 and edge device 5 has a defective link 26 to trunk switch 1.

When edge device 3 loses its connection to trunk switch 1 and maintains its connection 11 to trunk switch 2, then trunk switch 1 adds a transmit only connection 25 from the first switch means 13 to the intra cluster port 19 (double line arrow 25). Trunk switch 1 modifies those entries in address data base means 15 which direct packets to the port connected to edge device 3 to direct packets to the intra-cluster port 19. Also as previously described, trunk switch 2 adds a transmit only connection from the second switch means 27 to the edge port 9 connected to the edge device 3 (triple line arrow 33).

In addition if edge device 5 loses its connection 46 to trunk switch 1, and maintains its connection to trunk switch 2, then trunk switch 1 modifies those entries in address data base means 15 which direct packets to the port connected to edge device 3 to direct packets to the intra-cluster port 19. Trunk switch 2 adds a transmit only connection 44 from the second switch means 27 to the edge port 9 connected to the edge device 5 (triple line arrow 44).

In this way, individual trunk switches 1 and 2 of a trunk cluster are able to compensate for lost connections without needing to remove an entire trunk switch from the trunk cluster, and lose its corresponding bandwidth. The intra cluster data transmitted by intra cluster ports 19 and 29 is in addition to trunk synchronization data which is transferred between trunk switches to synchronize the address database and have all the trunk switches operate as a trunk cluster. The intra cluster ports 19 and 29, as well as the intra cluster link 35 is a high speed intra cluster interconnect (HSICI) and preferably operates in a manner substantially similar to the edge ports 9 and the links 22. The speeds of the ports 19 and 29, as well as the intra cluster link 35 is similar to the speeds of the edge ports 9 and the links 22.

FIG. 5 illustrates the traffic flow 40 through a multiple link trunked port of an edge device. FIG. 5 shows transmit queue 41, receive queue 42, traffic steering means 44, steering data base 46, media access controllers (MACs) 48, physical layers 49 and connected cable 22 (link 11). The trunked port 9 acts as a single port, but utilizes multiple physical links 22. The traffic steering means is the logic that knows which MAC 48 to use when transmitting a packet. The traffic steering means 44 dequeues transmit packets, examines the packet, consults the steering data base 46 and transfers the packet to the appropriate MAC 48. The steering data base 46 is constructed by information obtained by the reachability protocol or trunk configuration management protocol (TCMP).

Each trunk switch 1, 2 analyzes each received packet for the source address of the packet and the port on which the packet was received. This information is used to update the address data base means of the respective trunk switch, which is then used to update the address data base means of the other trunk switches. In this way, each trunk switch has not only knowledge of the identity of the edge devices, but knowledge of the addresses of further devices connected to the edge devices, such as individual workstations. The relationships of the further devices to the edge devices and to the edge ports, as detected by one trunk switch, is shared with the other trunk switches.

A trunk arrangement can have more than 2 trunk switches. In such an embodiment, at least one of the trunk switches would have more than one intra-cluster port 19. A trunk switch with two intra-cluster ports can connect to two other trunk switches. If a connection to an edge device is lost on a trunk switch with two intra-cluster ports, the trunk switch can send packets for that edge device to either of the two other trunk switches, providing their connections are operational.

Further details regarding operation of each trunk switch 1 and 2 with active links, ports, address data bases and edge devices can be found in Applicant's previously mentioned copending applications which have been incorporated by reference.

While specific embodiments of the invention have been shown and described in detail to illustrate the application of the principles of the invention, it will be understood that the invention may be embodied otherwise without departing from such principles.

What is claimed is:

1. A trunk switch arrangement comprising:

- a plurality of edge devices, each of said edge devices including a plurality of trunk links;
- a plurality of trunk switches, each of said trunk switches having a plurality of edge ports and being each separately connected to each one of said trunk links of each of said edge devices, said trunk switches cooperating to form a single logical entity receiving data from said edge devices and forwarding data to said edge devices, each edge device being able to communicate with all other edge devices through any one of said trunk switches;
- each of said trunk switches including an address data base for indicating a correspondence of said edge ports to the edge devices;
- each of said trunk switches including first switch means for reading data from each of said respective edge ports, consulting said respective address data base to determine which said respective edge port to send the received data out on, and transferring the received data to said determined respective edge port;

7

each of said trunk switches including an intra-cluster port connecting each of said trunk switches to another of said trunk switches forming the single logical entity; an intra-cluster link for connecting each intra-cluster port of one of said trunk switches to an intra-cluster port of another of said trunk switches and for transferring data between two said intra-cluster ports, each address data base means including synchronizing means for synchronizing address information with other said trunk switch means via said intra-cluster link;

said trunk switches including second switch means for reading intra-cluster data from a respective said intra-cluster port, consulting said respective address data base to determine which respective said edge port to send said intra-cluster data out on, and transferring said intra-cluster data to said determined respective edge port;

each of said trunk switches including control means for detecting a respective disconnected edge device, said control means modifying said respective address data base to direct data for said disconnected edge device to said respective intra-cluster port for transmission to another of said trunk switches;

said edge ports receive data in a form of a data packet with a destination address;

said first switch means transfers said data packet from one said respective edge port to one of another said respective edge ports and said intra-cluster port dependent on information present in said destination address and said address base means;

said intra-cluster port receives data in said form of data packets with a destination address;

said second switch means transfers said data packet from said respective intra-cluster port to one of said respective edge ports dependent on information present in said destination address and said address base means.

2. A trunk arrangement in accordance with claim 1, wherein:

said edge ports and said intra-cluster port transfer said data packets substantially identically.

3. A trunk arrangement in accordance with claim 1 wherein:

said edge ports and said intra-cluster port transfer said data at substantially identical speeds.

4. A trunk arrangement in accordance with claim 1, wherein:

said first switch means creates a transmit only connection from said first switch means to said intra-cluster port when said control means detects said disconnected edge device; and

said second switch means creates a transmit only connection from said second switch means to a disconnected edge device of another trunk switch of the trunk cluster.

5. A trunk arrangement in accordance with claim 1, wherein:

one of said trunk switch means includes another said intra-cluster port for connecting said one trunk switch means to another of said trunk switch means through another intra-cluster link;

said second switch means of said one trunk switch means reading further intra-cluster data from said another intra-cluster port, consulting said address data base means to determine which said edge port of said one trunk switch means to send said further intra-cluster data out on, and transferring said further intra-cluster data to said determined edge port;

8

said control means of said one trunk switch means modifying said address data base means of said one trunk switch means to redirect data for said disconnected edge device to one of said intra-cluster port and said another intra-cluster port for transmission to said other trunk switch means.

6. A trunk arrangement in accordance with claim 1, wherein:

each said address data base means includes correspondence between respective said edge ports and addresses of further devices connected to said edge devices.

7. A trunk arrangement in accordance with claim 6, wherein said trunk synchronous means synchronizes said correspondence between respective said edge ports and said addresses of further devices in said address database means of said trunk switches.

8. A method for trunking a plurality of edge devices through a plurality of trunk switches, the method comprising the steps of:

connecting each of the edge devices to separate edge port of each of the trunk switches to form a single logical switch with each of said trunk switches connected to each edge device;

providing an intra-cluster port on each of the trunk switches;

connecting each of the trunk switches to at least one other trunk switch through said intra-cluster ports;

constructing a steering data base using a reachability protocol or a trunk configuration management protocol; transmitting data from each edge device to said single logical switch with said transmitting of data being in a form of a data packet with a destination address;

determining at each trunk switch which edge port to send the received data out on, and transferring the received data to the determined edge port with said transferring the received data being dependent on information present in said destination address and said steering data base and a correspondence of said edge ports to the edge devices;

each trunk switch detecting when one of the edge devices has become disconnected from a respective trunk switch and updating said steering data base to change a correspondence between respective said edge ports and destination addresses to include one of said intra-cluster ports for one or more destination addresses;

synchronizing a correspondence between respective said edge ports and destination addresses to include one of said intra-cluster ports for one or more destination addresses by sending updated steering data from one steering data base to another steering data base via said intra-cluster link;

redirecting data at one of the trunk switches for the disconnected edge device to said intra-cluster port based on the updated steering data base;

based on the updated steering data base, each trunk switch reading intra-cluster data from a respective said intra-cluster port, determining which said respective edge port corresponds to said disconnected edge device of another trunk switch and transferring said intra-cluster data to said respective edge port corresponding to said disconnected edge device of the other trunk switch.

9. A method in accordance with claim 8, wherein:

said edge ports and said intra-cluster port transfer said data substantially identically.

9

10. A method in accordance with claim 8, wherein:

said edge ports and said intra-cluster port transfer said data at substantially identical speeds.

11. A method in accordance with claim 8, further comprising:

providing another intra-cluster port on one of the trunk switches;

connecting each of the trunk switches to each other through said intra-cluster ports and said another intra-cluster ports;

said one trunk switch detecting when one of the edge devices has become disconnected from said one trunk switch and redirecting data for the one disconnected edge device to one of said intra-cluster port and said another intra-cluster port;

said one trunk switch reading further intra-cluster data from said another intra-cluster port, determining which said respective edge port corresponds to said disconnected edge device of another trunk switch and transferring said further intra-cluster data to said respective edge port corresponding to said disconnected edge device of the other trunk switch.

12. A method in accordance with claim 8, wherein:

each trunk switch determines correspondence between respective said edge ports and source addresses of further devices connected to said edge devices.

13. A logical switch with a plurality of trunk switches forming a single logical entity, the single logical entity comprising:

a plurality of edge ports for each of said trunk switches with each edge port connected to a trunk link connected to a different edge device;

an intra-cluster port for each of said trunk switches;

a logical entity internal path comprising at least a plurality of intra-cluster links connecting each intra-cluster port

10

of one of said trunk switches to an intra-cluster port of another said trunk switches and for transferring data between two said intra-cluster ports connecting each of said trunk switches to another of said trunk switches forming the single logical device;

an address data base for each of said trunk switches for indicating a correspondence of said edge ports to the edge devices or for indicating a correspondence of said intra-cluster ports to one or more edge devices upon one or more of the edge devices becoming disconnected from a respective trunk switch, each of said trunk switches reading data from data packets received at each of said respective edge ports, consulting said respective address data base to determine which respective edge port or intra-cluster port to send the received data out on, and transferring the received data to said determined respective edge port or intra-cluster port and each of said trunk switches reading data from data packets from a respective said intra-cluster port, consulting said respective address data base to determine which respective said edge port or intra-cluster port to send said intra-cluster data out on, and transferring said intra-cluster data to said determined respective edge port or intra-cluster port; and

a trunk switch control for each of said trunk switches for detecting a respective disconnected edge device, said control means modifying said respective address data base to direct data for said disconnected edge device to said respective intra-cluster port for transmission to another of said trunk switches and sending modifying data to each respective address data base of each other switch, via said intra cluster link to synchronize each address data base of each trunk switch of the single logical entity.

* * * * *